

A λ Prolog Based Animation of Twelf Specifications

Mary Southern and Gopalan Nadathur

University of Minnesota, Minneapolis MN 55455, USA

Abstract. Specifications in the Twelf system are based on a logic programming interpretation of the Edinburgh Logical Framework or LF. We consider an approach to animating such specifications using a λ Prolog implementation. This approach is based on a lossy translation of the dependently typed LF expressions into the simply typed lambda calculus (STLC) terms of λ Prolog and a subsequent encoding of lost dependency information in predicates that are defined by suitable clauses. To use this idea in an implementation of logic programming *a la* Twelf, it is also necessary to translate the results found for λ Prolog queries back into LF expressions. We describe such an inverse translation and show that it has the necessary properties to facilitate an emulation of Twelf behavior through our translation of LF specifications into λ Prolog programs. A characteristic of Twelf is that it permits queries to consist of types which have unspecified parts represented by meta-variables for which values are to be found through computation. We show that this capability can be supported within our translation based approach to animating Twelf specifications.

1 Introduction

The Edinburgh Logical Framework or LF [4] is a dependently typed lambda calculus that has proven useful in specifying formal systems such as logics and programming languages (see, e.g., [5]). The key to its successful application in this setting is twofold. First, the abstraction operator that is part of the syntax of LF provides a means for succinctly encoding formal objects whose structures embody binding notions. Second, LF types can be indexed by terms and, as such, they can be used to represent relations between objects that are encoded by terms. More precisely, types can be viewed as formulas and type checking as a means for determining if a given term represents a proof of that formula. Proof search can be introduced into this context by interpreting a type as a request to determine if there is a term of that type. Further, parts of a type can be left unspecified, thinking of it then as a request to fill in these parts in such a way that the resulting type is inhabited. Interpreting types in this way amounts to giving LF a logic programming interpretation. The Twelf system [9,10] is a realization of LF that is based on such an interpretation.

An alternative approach to specifying formal systems is to use a predicate logic. Objects treated by the formal systems can be represented by the terms of

this logic and relations between them can be expressed through predicates over these terms. If the terms include a notion of abstraction, e.g., if they encompass simply typed lambda terms, then they provide a convenient means for representing binding notions. By restricting the formulas that are used to model relations suitably, it is possible to constrain proof search behavior so that the formulas can be given a rule-based interpretation. The logic of higher-order hereditary Harrop formulas (*hohh*) has been designed with these ideas in mind and many experiments have shown this logic to be a useful specification device (see, e.g., [7]). This logic has also been given a computational interpretation in the language λ Prolog [8], for which efficient implementations such as the Prolog/Mali [1] and the Teyjus [11] systems have been developed.

The two different approaches to specification that are described above have a relationship that has been explored formally. In early work, Felty and Miller showed that LF derivations could be encoded in *hohh* derivations by describing a translation from the former to the latter [3]. This translation demonstrated the expressive power of *hohh*, but did not show the correspondence in proof search behavior. To rectify this situation, Snow et. al. described a transformation of LF specifications into *hohh* formulas that allowed the construction of derivations to be related [12]. This work also showed how to make the translation more efficient by utilizing information available from a static checking of LF types, and it refined the resulting *hohh* specifications towards making their structure more closely resemble that of the LF specifications they originated from.

The primary motivation for the work of Snow et. al. was a desire to use Teyjus as a backend for an alternative implementation of logic programming in Twelf. However, it falls short of achieving this goal in two ways that we address in this paper. First, although it relates derivations from LF specifications to ones from their translations, it does not make explicit the process of extracting an LF “result” term from a successful *hohh* derivation; such an extraction is necessary if Teyjus is to serve as a genuine, invisible backend. To close this gap, we describe an inverse translation and show that it has the necessary properties to allow Twelf behavior to be emulated through computations from λ Prolog programs. Second, Snow et. al. dealt only with closed types, i.e., they did not treat the idea of filling in missing parts of types in the course of looking for an inhabitant. To overcome this deficiency, we include meta-variables in specifications and treat them in the back-and-forth translations as well as in derivations; the last aspect, that is also the most critical one in our analysis, requires us to build substitutions and unification explicitly into our formalization of derivations.

The remainder of this paper is structured as follows. Sections 2 and 3 respectively present LF and the *hohh* logic together with their computational interpretations. Section 4 describes a translation from LF specifications into *hohh* ones together with an inverse translation for extracting solution terms from *hohh* derivations. We then propose an approach for developing a proof of correctness for this translation. Section 5 improves the basic translation and Section 6 uses it to illustrate our proposed approach to realizing logic programming in Twelf. Section 7 concludes the paper.

$$\begin{array}{c}
\frac{X : A \in \Delta}{\Gamma \vdash_{\Sigma} X : A} \text{ meta-var} \\
\\
\frac{\Sigma \text{ sig } c : A \in \Sigma}{\Gamma \vdash_{\Sigma} c : A^{\beta}} \text{ const-obj} \quad \frac{\Gamma \vdash_{\Sigma} A : \text{Type} \quad \Gamma, x : A \vdash_{\Sigma} M : B}{\Gamma \vdash_{\Sigma} (\lambda x : A. M) : (\Pi x : A^{\beta}. B)} \text{ abs-obj} \\
\\
\frac{\vdash_{\Sigma} \Gamma \text{ ctx } x : A \in \Gamma}{\Gamma \vdash_{\Sigma} x : A^{\beta}} \text{ var-obj} \quad \frac{\Gamma \vdash_{\Sigma} M : \Pi x : A. B \quad \Gamma \vdash_{\Sigma} N : A}{\Gamma \vdash_{\Sigma} (M N) : (B[N/x])^{\beta}} \text{ app-obj}
\end{array}$$

Fig. 1. Rules for typing LF objects

2 Logic programming in LF

Three categories of expressions constitute LF: kinds, type families or types which are classified by kinds, and objects which are classified by types. Below, x denotes an object variable, X an object meta-variable, c an object constant, and a a type constant. Letting K range over kinds, A and B over types, and M and N over objects, the syntax of these expressions is given as follows:

$$\begin{array}{lcl}
K & ::= & \text{Type} \mid \Pi x : A. K \\
A, B & ::= & a \mid \Pi x : A. B \mid A M \\
M, N & ::= & c \mid x \mid X \mid \lambda x : A. M \mid M N
\end{array}$$

Both Π and λ are binders which also assign types to the (object) variables they bind over expressions. Notice the dependency present in LF expressions: a bound object variable may appear in a type family or kind. In the remainder of this paper we use U and V ambiguously for types and objects and P similarly for types and kinds. The shorthand $A \rightarrow P$ is used for $\Pi x : A. P$ if P is a type family or kind that is not dependent on the bound variable, i.e. if x does not appear free in P . Terms differing only in bound variable names are identified. We write $U[M_1/x_1, \dots, M_n/x_n]$ to denote the capture avoiding substitution of M_1, \dots, M_n for the free occurrences of x_1, \dots, x_n respectively in U .

LF kinds, types and objects are formed relative to a signature Σ that identifies constants together with their kinds or types. In determining if an expression is well-formed, we additionally need to consider contexts, denoted by Γ , that assign types to variables. The syntax for signatures and contexts is as follows:

$$\Sigma ::= \cdot \mid \Sigma, a : K \mid \Sigma, c : A \quad \Gamma ::= \cdot \mid \Gamma, x : A$$

In contrast to usual LF presentations, we have allowed expressions to contain object meta-variables. We assume an infinite supply of such variables for each type and that an implicit meta-variable context Δ assigns types to these variables. These meta-variables act as placeholders, representing the part of an expression one wishes to leave unspecified.

Complementing the syntax rules, LF has typing rules that limit the set of acceptable or well-formed expressions. These rules define the following mutually recursive judgments with the associated declarative content:

$$\Sigma \text{ sig} \quad \Sigma \text{ is a valid signature}$$

$\vdash_{\Sigma} \Gamma \text{ ctx}$	Γ is a valid context relative to the (valid) signature Σ
$\Gamma \vdash_{\Sigma} K \text{ kind}$	K is a valid kind in signature Σ and context Γ
$\Gamma \vdash_{\Sigma} A : K$	A is a type of kind K in a signature Σ and context Γ
$\Gamma \vdash_{\Sigma} M : A$	M is an object of type A in signature Σ and context Γ

In our discussion of logic programming, we rely on a specific knowledge of the rules for only the last of these judgments which we present in Figure 1; an intuition for the other rules should follow from the ones presented and their explicit presentation can be found, e.g., in [4]. By these rules we can see that if a well-formed expression contains a meta- variable X of type A , then replacing the occurrences of X with a well- formed object of type A will produce an expression which is also well-formed.

The rules in Figure 1 make use of an equality notion for LF expressions that is based on β -conversion, i.e., the reflexive and transitive closure of a relation equating two expressions which differ only in that a subexpression of the form $((\lambda x:A.M) N)$ in one is replaced by $M[N/x]$ in the other. We shall write U^{β} for the β -normal form of an expression, i.e., for an expression that is equal to U and that does not contain any subexpressions of the form $((\lambda x:A.M) N)$. Such forms are not guaranteed to exist for all LF expressions. However, they do exist for well-formed LF expressions [4], a property that is ensured to hold for each relevant LF expression by the premises of every rule whose conclusion requires the β -normal form of that expression.

Equality for LF expressions also includes η -conversion, i.e., the congruence generated by the relation that equates $\lambda x:A.(M x)$ and M if x does not appear free in M . The β -normal forms for the different categories of expressions have the following structure

$$\begin{array}{ll}
\textit{Kind} & \Pi x_1:A_1 \dots \Pi x_n:A_n. \textit{Type} \\
\textit{Type} & \Pi y_1:B_1 \dots \Pi y_m:B_m. a \ M_1 \ \dots \ M_n \\
\textit{Object} & \lambda x_1:A_1 \dots \lambda x_n:A_n. u \ M_1 \ \dots \ M_n
\end{array}$$

where u is an object constant or variable and where the subterms and subtypes appearing in the expression recursively have the same form. We refer to the part corresponding to $a \ M_1 \ \dots \ M_n$ in a type in this form as its *target* type and to B_1, \dots, B_m as its *argument* types. Let w be a variable or constant which appears in the well-formed term U and let the number of Π s that appear in the prefix of its type or kind in beta normal form be n . We say w is *fully applied* if every occurrence of w in U has the form $w \ M_1 \ \dots \ M_n$. A type of the form $a \ M_1 \ \dots \ M_n$ where a is fully applied is a *base type*. We also say that U is *canonical* if it is in normal form and every occurrence of a variable or constant in it is fully applied. It is a known fact that every well-formed LF expression is equal to one in canonical form by virtue of $\beta\eta$ -conversion [4]. For the remainder of this paper we will assume all terms are in β -normal form.

A specification in LF comprises a signature that, as we have seen, identifies a collection of object and type constants. The Curry-Howard isomorphism [6] allows types to be interpreted dually as formulas. The dependent nature of the

```

nat : type.                list : type.
z : nat.                   nil : list.
s : nat -> nat.            cons : nat -> list -> list.

append : list -> list -> list -> type.
app-nil : append nil L L.
app-cons : append L1 L2 L3 -> append (cons X L1) L2 (cons X L3).

```

Fig. 2. A Twelf signature specifying lists and the append relation

LF type system allows type constants to take objects as arguments. Such constants then correspond to the names of predicates over suitably typed objects. Moreover, the same isomorphism allows object constants, which provide a means for constructing expressions of particular types, to be viewed as the names of parameterized rules for constructing proofs of the relations represented by the types.

Figure 2 presents a concrete signature to illustrate these ideas. In showing this and other similar signatures, we use the Twelf syntax for LF expressions. In this syntax, $\Pi x:A.U$ is written as $\{x : A\} U$ and $\lambda x:A.M$ is written as $[x : A] M$. Further, bindings and the corresponding type annotations on variables are made implicit in situations where the types can be uniquely inferred; the variables that are implicitly bound are denoted in Prolog style by tokens that begin with uppercase letters. The initial part of the signature in Figure 2 defines type and object constants that provide a representation of the natural numbers and lists of natural numbers. The signature then identifies a type constant **append** that takes three lists as arguments. Under the viewpoint just explained, this constant can be interpreted as a predicate that relates three lists. Objects of this type can be constructed by using the constants **app-nil** and **app-cons** that are also presented in the signature. Viewed differently, these constants name rules that can be used to construct a proof of the append relation between three lists. Notice that **app-cons** requires as an argument an object of **append** type. This object plays the role of a premise for the rule that **app-cons** identifies.

The logic programming use of LF that underlies Twelf consists of presenting a type A in the setting of a signature Σ . Such a type corresponds to the request to find an object M such that the judgment $\vdash_{\Sigma} M : A$ is derivable. Alternately, a query in Twelf can be seen as the desire to determine the derivability of a formula, the inhabiting term that is found being its proof. The type that is presented as a query may also contain meta-variables, denoted by tokens that begin with uppercase letters. In this case, the request is to find substitutions for these variables while simultaneously showing that the instance type is inhabited.

An example of a query relative to the signature in Figure 2 is the following.

```
append (cons z nil) nil L
```

An answer to this query is the substitution $(\text{cons } z \text{ nil})$ for L , together with the object $(\text{app-cons } (\text{cons } z \text{ nil}) \text{ nil } (\text{cons } z \text{ nil}) (\text{app-nil nil}))$ that inhabits that type. Another query in this setting is

```
{x:nat} append (cons x nil) (cons z (cons x nil)) (L x).
```

$$\begin{array}{c}
\frac{}{\Xi; \Gamma \longrightarrow \top} \top R \quad \frac{\Xi; \Gamma \cup \{D\} \longrightarrow G}{\Xi; \Gamma \longrightarrow D \supset G} \supset R \quad \frac{c \notin \Xi \quad \Xi \cup \{c\}; \Gamma \longrightarrow G[c/x]}{\Xi; \Gamma \longrightarrow \forall x. G} \forall R \\
\\
\frac{\Xi; \Gamma \longrightarrow G_1[\overrightarrow{t_1/x_1}] \quad \dots \quad \Xi; \Gamma \longrightarrow G_n[\overrightarrow{t_1/x_1}, \dots, \overrightarrow{t_n/x_n}]}{\Xi; \Gamma \longrightarrow A} \text{backchain} \\
\text{where } \forall \overrightarrow{x_1}. (G_1 \supset \dots \supset \forall \overrightarrow{x_n}. (G_n \supset A')) \dots \in \Gamma, \\
\overrightarrow{t_1}, \dots, \overrightarrow{t_n} \text{ are } \Xi\text{-terms and } A'[\overrightarrow{t_1/x_1}, \dots, \overrightarrow{t_n/x_n}] = A
\end{array}$$

Fig. 3. Derivation rules for the *hohh* logic

in which L is a “higher-order” meta-variable of type `nat -> list`. The substitution that would be computed by Twelf for the variable L in this query is

`[y:nat] (cons y (cons z (cons y nil)))`,

and the corresponding inhabitant or proof term is

`[y:nat] app-cons nil (cons z (cons y nil))`
`(cons z (cons y nil)) y`
`(app-nil (cons z (cons y nil)))`

Notice that the variable x that is explicitly bound in the query has a *different* interpretation from the meta-variable L . In particular, it receives a “universal” reading: the query represents a request to find a value for L that yields an inhabited type regardless of what the value of x is.

Although neither of our example queries exhibited this behavior, the range of an answer substitution may itself contain variables and there may be some residual constraints on these variables presented in the form of a collection of equations between object expressions called “disagreement pairs.” The interpretation of such an answer is that a complete solution can be obtained from the provided substitution by instantiating the remaining variables with closed object expressions that render identical the two sides of each disagreement pair.

3 Logic programming based on *hohh*

An alternative approach to specifying formal systems is to use a logic in which relationships between terms are encoded in predicates. The idea of animating a specification then corresponds to constructing a proof for a given “goal” formula in the chosen logic. To yield a sensible notion of computation, specifications must also be able to convey information about how a search for a proof should be conducted. Towards this end, we use here the logic of higher-order hereditary Harrop formulas, referred to in short as the *hohh* logic. This logic underlies the programming language λ Prolog [8].

The *hohh* logic is based on Church’s Simple Theory of Types [2]. The expressions of this logic are those of a simply typed λ -calculus (STLC). Types are constructed from the atomic type o for propositions and a finite set of other atomic types by using the function type constructor \rightarrow . We assume we have been

```

nat : type.          list : type.
z : nat.             nil : list.
s : nat -> nat.       cons : nat -> list -> list.
                     append : list -> list -> list -> o.

∀L. append nil L L.
∀X∀L1∀L2∀L3. append L1 L2 L3 ⊃ append (cons X L1) L2 (cons X L3).

```

Fig. 4. An *hohh* specification of lists and the append relation

given a set of variables and a set of constants, each member of these sets being identified together with a type. More complex terms are constructed from these atomic symbols by using application and λ -abstraction in a way that respects the constraints of typing. As in LF, terms differing only in bound variable names are identified. The notion of equality between terms is further enriched by β - and η -conversion. When we orient these rules and think of them as reductions, we are assured in the simply typed setting of the existence of a unique normal form for every well-formed term under these reductions. Thus, equality between two terms becomes the same as the identity of their normal forms. For simplicity, in the remainder of this paper we will assume that all terms have been converted to normal form. We write $t[s_1/x_1, \dots, s_n/x_n]$ to denote the capture avoiding substitution of the terms s_1, \dots, s_n for free occurrences of x_1, \dots, x_n in t .

Logic is introduced into this setting by identifying a sub-collection of the set of constants as logical constants and giving them a special meaning. The logical constants that we shall use here are the following:

\top of type o
 \supset of type $o \rightarrow o \rightarrow o$
 Π of type $(\tau \rightarrow o) \rightarrow o$ for each type τ

We intend \top to denote the always true proposition and \supset , which we will write in infix form, to denote implication. The symbol Π corresponds to the generalized universal quantifier: the usual notation $\forall x.F$ for universal quantification serves as a shorthand for $\Pi(\lambda x.F)$.

To construct a specification within the *hohh* logic, a user must identify a collection of types and a further set of constants, called non-logical constants, together with their types. A collection of such associations forms a signature. There is a proviso on the types of non-logical constants: their argument types must not contain o . Non-logical constants that have o as their target or result type correspond to predicate symbols. If c is such a constant with the type $\tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow o$ and t_1, \dots, t_n are terms of type τ_1, \dots, τ_n , respectively, then the term $(c \ t_1 \ \dots \ t_n)$ of type o constitutes an *atomic formula*. We shall use the syntax variable A to denote such formulas. More complex terms of type o are constructed from atomic formulas by using the logical constants. Such terms are also referred to as *formulas*.

The *hohh* logic is based on two special classes of formulas identified by the following syntax rules:

$$G ::= \top \mid A \mid D \supset G \mid \forall x.G \qquad D ::= A \mid G \supset D \mid \forall x.D$$

$$\begin{array}{ll}
\phi(A) := \text{lf-obj} \text{ when } A \text{ is a base type} & \langle u \rangle := u \\
\phi(\Pi x:A.P) := \phi(A) \rightarrow \phi(P) & \langle x \rangle := x \\
\phi(\text{Type}) := \text{lf-type} & \langle X \rangle := X \\
& \langle M_1 \ M_2 \rangle := \langle M_1 \rangle \ \langle M_2 \rangle \\
& \langle \lambda x:A.M \rangle := \lambda^{\phi(A)} x. \langle M \rangle
\end{array}$$

Fig. 5. Flattening of types and encoding of terms

We will refer to a D -formula also as a program clause. Notice that, in elaborated form, such a formula has the structure $\forall \vec{x}_1. (G_1 \supset \dots \supset \forall \vec{x}_n. (G_n \supset A) \dots)$; we write $\forall \vec{x}_i$ here to denote a sequence of universal quantifications.

The computational interpretation of the *hohh* logic consists of thinking of a collection of D -formulas as a program and a G -formula as a goal or query that is to be solved against a given program \mathcal{P} in the context of a given signature Ξ . We represent the judgment that the query G has a solution in such a setting by the “sequent” $\Xi; \mathcal{P} \longrightarrow G$. The rules for deriving such a judgment are shown in Figure 3. Using these rules to search for a derivation leads to a process in which we first simplify a goal in a manner determined by the logical constants that appear in it and then employ program clauses in a familiar backchaining mode to solve the atomic goals that are produced. A property of the *hohh* logic that should be noted is that both the program and the signature can change in the course of a computation.

We illustrate the use of these ideas in practice by considering, once again, the encoding of lists of natural numbers and the append relation on them. Figure 4 provides both the signature and the program clauses that are needed for this purpose. This specification is similar to one that might be provided in Prolog, except for the use of a curried notation for applications and the fact that the language is now typed. We “execute” these specifications by providing a goal formula. As with Twelf, we will allow goal formulas to contain free or meta-variables for which we intend instantiations to be found through proof search. A concrete example of such a goal relative to the specification in Figure 4 is `(append (cons z nil) nil L)`. This goal is solvable with the substitution `(cons z nil)` for `L`. Another example of a query in this setting is $\forall x. (\text{append} (\text{cons } x \text{ nil}) (\text{cons } z (\text{cons } x \text{ nil})) (L \ x))$ and an answer to this goal is the substitution $\lambda y. (\text{cons } y (\text{cons } z (\text{cons } y \text{ nil})))$ for `L`.

4 Translating Twelf specifications into predicate form

We now turn to the task of animating Twelf specifications using a λ Prolog implementation. Towards this end, we describe a meaning preserving translation from LF signatures into *hohh* specifications. Our translation extends the one in [12] by allowing for meta-variables in LF expressions. We also present an inverse translation for bringing solutions back from λ Prolog to the Twelf setting.

$$\begin{aligned}\llbracket \Pi x:A.B \rrbracket &:= \lambda M. \forall x. (\llbracket A \rrbracket x) \supset (\llbracket B \rrbracket (M x)) \\ \llbracket A \rrbracket &:= \lambda M. \text{hastype } M \langle A \rangle \text{ where } A \text{ is a base type}\end{aligned}$$

Fig. 6. Encoding of LF types using the **hastype** predicate

The first step in our translation is to map dependently typed lambda expressions into simply typed ones. We shall represent both types and objects in LF by STLC terms (which are also *hohh* terms), differentiating the two categories by using the (simple) type *lf-obj* for the encodings of LF objects and *lf-type* for those of LF types. To play this out in detail, we first associate an *hohh* type with each LF type and kind that is given by the $\phi(\cdot)$ mapping shown in Figure 5. Then, corresponding to each object and type-level LF constant $u : P$, we identify an *hohh* constant with the same name but with type $\phi(P)$. Finally, we transform LF objects and kinds into *hohh* terms using the $\langle \cdot \rangle$ mapping in Figure 5.

We would like to consider an inverse to the transformation that we have described above. We have some extra information available in constructing such an inverse: the constants that appear in the *hohh* terms of interest have their correlates which have been given specific types in the originating LF signature. Even so, the lossy nature of the translation makes the inversion questionable. There are two kinds of problems. First, because (the chosen) simple typing is not sufficiently constraining, we may have well-formed STLC terms for which there is no corresponding LF expression. As a concrete example, consider the following LF signature:

$i : \text{type} \quad j : \text{type} \quad a : i \rightarrow j \quad c : i$

In the encoding we will have the following two constants with associated types:

$a : \text{lf-obj} \rightarrow \text{lf-obj} \quad c : \text{lf-obj}$

This means that we can construct the simply typed term $(a (a c))$ which cannot be the image of any LF expression that is well-formed under the given signature. The second problem is that when an *hohh* term involves an abstraction, the choice of LF type to use for the abstracted variable is ambiguous. As a concrete example, consider the *hohh* term $\lambda x.x$ that has the type $\text{lf-obj} \rightarrow \text{lf-obj}$. This term could map to the LF objects $[x:\text{nat}] x$ and $[x:\text{list}] x$, amongst many other choices.

Our solution to these problems is twofold. First, we will assume that we know the type of the LF object that the inversion is to produce; this information will always be available when the *hohh* terms arise in the course of simulating LF typing derivations using *hohh* derivations. Second, we will define inversion as a partial function: when we use it to calculate an LF expression from an answer substitution returned by an *hohh* computation, we will have an additional obligation to show that the inverse must exist.

The rules in Figure 7 define the inverse transformation. The judgments $\text{inv}^\downarrow(t; A; \Theta) = M$ and $\text{inv}^\uparrow(t; A; \Theta) = M$ are to be derivable when t is an

$$\begin{array}{c}
\frac{X : A \in \Delta}{inv^\uparrow(X; A; \Theta) = X} \text{ inv-var} \qquad \frac{inv^\downarrow(M; B; \Theta, x : A) = M'}{inv^\downarrow(\lambda x.M; \Pi x:A.B; \Theta) = \lambda x:A.M'} \text{ inv-abs} \\
\\
\frac{inv^\uparrow(M_1; \Pi x:B.A; \Theta) = M'_1 \quad inv^\downarrow(M_2; B; \Theta) = M'_2}{inv^\uparrow(M_1 \ M_2; A[M'_2/x]; \Theta) = M'_1 \ M'_2} \text{ inv-app} \\
\\
\frac{u : A \in \Theta}{inv^\uparrow(u; A; \Theta) = u} \text{ inv-const} \qquad \frac{inv^\uparrow(M; A; \Theta) = M'}{inv^\downarrow(M; A; \Theta) = M'} \text{ inv-syn}
\end{array}$$

Fig. 7. An inverse encoding

hohh term in β -normal form that inverts to the LF object M that has type A in a setting where variables and constants are typed according to Θ . The difference between the two judgments is that the first expects A as an input whereas the second additionally synthesizes the type. The process starts with checking against an LF type—this type will be available from the original LF query—and it is easily shown that if $inv^\downarrow(t; A; \Sigma \cup \Gamma) = M$, then $\Gamma \vdash_\Sigma M : A$. Notice that we will only ever check an abstraction term against an LF type, ensuring that the type chosen for the bound variable will be unique. We say a substitution θ is invertible in a given context and signature if each term in its range is invertible in that setting, using the type associated with the domain variable by Δ .

The translation of LF expressions into *hohh* terms loses all relational information encoded by dependencies in types. For example it transforms the constants encoding the append relation in Figure 2 into the following *hohh* signature:

```

append : lf-obj -> lf-obj -> lf-obj -> lf-type.
app-nil : lf-obj -> lf-obj.
app-cons : lf-obj -> lf-obj ->
           lf-obj -> lf-obj -> lf-obj -> lf-obj.

```

It is no longer possible to construe this as a specification of the append relation between lists. To recover the lost information, we employ a second pass that uses predicates to encode relational content. This pass employs the *hohh* predicate *hastype* with type $lf-obj \rightarrow lf-type \rightarrow o$ and generates clauses that are such that *hastype* $X \ T$ is derivable from them exactly when X is the encoding of an LF term M of a base LF type whose encoding is T . More specifically, this pass processes each item of the form $U : P$ in the LF signature and produces from it the clause $\{\{P\}\} \langle U \rangle$ using the rules in Figure 6 that define $\{\{.\}\}$.

To illustrate the second pass, when used with the signature in Figure 2, we see that it will produce the following clauses:

```

hastype z nat.
∀x.hastype x nat ⊃ hastype (s x) nat.
hastype nil list.
∀x.(hastype x nat ⊃
    ∀l.(hastype l list ⊃ hastype (cons x l) list)).

∀l.hastype l list ⊃ hastype (app-nil l) list.

```

$$\begin{aligned} &\forall x. (\text{hastype } x \text{ nat} \supset \forall l1. (\text{hastype } l1 \text{ list} \supset \\ &\quad \forall l2. (\text{hastype } l2 \text{ list} \supset \forall l3. (\text{hastype } l3 \text{ list} \supset \\ &\quad \quad \forall a. (\text{hastype } a \text{ (append } l1 \text{ } l2 \text{ } l3)} \supset \\ &\quad \quad \quad \text{hastype (app-cons } x \text{ } l1 \text{ } l2 \text{ } l3 \text{ } a) \\ &\quad \quad \quad \quad (\text{append (cons } x \text{ } l1) \text{ } l2 \text{ (cons } x \text{ } l3)))))). \end{aligned}$$

Contrasting these clauses with the ones of the λ Prolog program in Figure 4, we see that it is capable not only of producing answers to `append` queries but also a “proof-term” that traces the derivation of such queries.

The correctness of our translation is captured by the following theorem (whose proof is currently incomplete). We had said earlier that when looking at terms that are produced by *hohh* derivations from LF translations, we would have an assurance that these terms are invertible. This is a property that flows, in fact, from the structure of the *hastype* clauses: as a *hohh* derivation is constructed, all the substitution terms that are generated are checked to be of the right type using the *hastype* predicate, and so we will not be able to construct a term which is not invertible.

Theorem 1. *Let Σ be an LF signature and let A be an LF type that possibly contains meta-variables.*

1. *If Twelf solves the query $M : A$ with the ground answer substitution σ , then there is an invertible answer substitution θ for the goal $\{\!\{A\}\!\} \langle M \rangle$ wrt $\{\!\{\Sigma\}\!\}$ such that the inverse θ' of θ generalizes σ (i.e. there exists a σ' such that $\sigma' \circ \theta' = \sigma$).*
2. *If θ is an invertible answer substitution for $\{\!\{A\}\!\} \langle M \rangle$, then its inverse is an answer substitution for $M : A$.*

Our approach to proving this theorem is to consider the operational semantics of the two systems and to show that derivations in each system can be factored into sequences of steps that can be simulated by the other system. Moreover, this simulation ensures the necessary relationships hold between the answer substitutions that are gradually developed by the derivations in the respective systems.

5 Optimizing the translation

The translation presented in the preceding section does not lend itself well to proof search because it generates a large amount of redundant typing checking. There are many instances when this redundancy can be recognized by a direct analysis of a given Twelf specification: in particular, we can use a structural analysis of an LF expression to determine that a term being substituted for a variable must be of the correct type and hence it is unnecessary to check this explicitly. In this section we develop this idea and present an improved translation. We also discuss another optimization that reflect the types in the Twelf signature more directly into types in *hohh*. The combination of these optimizations produce clauses that are more compact and that resemble those that might be written in λ Prolog directly.

$$\begin{array}{c}
\frac{\text{dom}(\Gamma); \cdot; x \sqsubseteq_o A_i \text{ for some } A_i \text{ in } \vec{A}}{\Gamma; x \sqsubseteq_t c \vec{A}} \text{APP}_t \\
\frac{\Gamma, y : A; x \sqsubseteq_t B}{\Gamma; x \sqsubseteq_t \Pi y : A.B} \text{PI}_t \quad \frac{\Gamma_1; x \sqsubseteq_t B \quad \Gamma_1, y : B, \Gamma_2; y \sqsubseteq_t A}{\Gamma_1, y : B, \Gamma_2; x \sqsubseteq_t A} \text{CTX}_t \\
\frac{y_i \in \delta \text{ for each } y_i \text{ in } \vec{y} \quad \text{each variable in } \vec{y} \text{ is distinct}}{\Delta; \delta; x \sqsubseteq_o x \vec{y}} \text{INIT}_o \\
\frac{y \notin \Delta \text{ and } \Delta; \delta; x \sqsubseteq_o M_i \text{ for some } M_i \text{ in } \vec{M}}{\Delta; \delta; x \sqsubseteq_o y \vec{M}} \text{APP}_o \quad \frac{\Delta; \delta; y; x \sqsubseteq_o M}{\Delta; \delta; x \sqsubseteq_o \lambda y : A.M} \text{ABS}_o
\end{array}$$

Fig. 8. Strictly occurring variables in types and objects

We are interested in translating an LF type of the form $\Pi x_1 : A_1 \dots \Pi x_n : A_n.B$ into an *hohh* clause that can be used to determine if a type B' can be viewed as an instance $B[M_1/x_1, \dots, M_n/x_n]$ of the target type B . This task also requires us to show that M_1, \dots, M_n are inhabitants of the types A_1, \dots, A_n ; in the naive translation, this job is done by the *hastype* formulas pertaining to x_i and A_i that appear in the body of the *hohh* clause produced for the overall type. However, a particular x_i may occur in B in a manner which already makes it clear that the term M_i which replaces it in any instance of B must possess such a property. What we want to do, then, is characterize such occurrences of x_i such that we can avoid having to include an inhabitation check in the *hohh* clause.

We define a strictness condition for variable occurrences and, hence, for variables that possesses this kind of property. By using this condition, we can simplify the translation of a type into an *hohh* clause without losing accuracy. In addition to efficiency, such a translation also produces a result that bears a much closer resemblance to the LF type from which it originates.

The critical idea behind this criterion is that the path down to the occurrence of x is *rigid*, i.e., it cannot be modified by substitution and x is not applied to arguments in a way that could change the structure of the expression substituted for it. We know that the structure will be unchanged by application of arguments by requiring the occurrence of x to be applied only to distinct λ -bound variables. Thus we know that any term substituted for x has the correct type without needing to explicitly check it. Specifically, we say that the bound variable x_i occurs strictly in the type $\Pi x_1 : A_1 \dots \Pi x_n : A_n.B$ if it is the case that

$$x_1 : A_1, \dots, x_{i-1} : A_{i-1}; x_i \sqsubseteq_t \Pi x_{i+1} : A_{i+1} \dots \Pi x_n : A_n.B$$

holds. We have been able to extend the strictness condition as described in [12] recursively while preserving its utility in recognizing redundancy in type checking. We consider occurrences of bound variables to be strict in the overall type if they are strict in the types of other bound variables that occur strictly in the target type. The relation defined in Figure 8 formalizes this idea.

When $\Gamma; x \sqsubseteq_t A$ is derivable it means that the variable x appears strictly in the type A in the context Γ . As we work down through the structure of a type we will eventually look at a specific term M and a derivation of $\Delta; \delta; x \sqsubseteq_o M$ means that x appears strictly in the term M . Here, Δ and δ are both lists of

$$\begin{aligned}
\phi(a \ M_1 \dots M_n) &:= a\text{-type} & \langle u \rangle &:= u \\
\phi(\Pi x:A.P) &:= \phi(A) \rightarrow \phi(P) & \langle x \rangle &:= x \\
\phi(\text{Type}) &:= \text{lf-type} & \langle X \rangle &:= X \\
& & \langle M_1 \ M_2 \rangle &:= \langle M_1 \rangle \ \langle M_2 \rangle \\
& & \langle \lambda x:A.M \rangle &:= \lambda^{\phi(A)} x. \langle M \rangle
\end{aligned}$$

$$\begin{aligned}
\llbracket \Pi x:A.B \rrbracket_F^+ &:= \begin{cases} \lambda M. \forall x. \top \supset \llbracket B \rrbracket_{F,x}^+(M \ x) & \text{if } \Gamma; x \sqsubset_t B \\ \lambda M. \forall x. \llbracket A \rrbracket^-(x) \supset \llbracket B \rrbracket_{F,x}^+(M \ x) & \text{otherwise} \end{cases} \\
\llbracket u \ \vec{N} \rrbracket_F^+ &:= \lambda M. u \ \overrightarrow{\langle N \rangle} M \\
\llbracket \Pi x:A.B \rrbracket^- &:= \lambda M. \forall x. \llbracket A \rrbracket^+(x) \supset \llbracket B \rrbracket^-(M \ x) \\
\llbracket u \ \vec{N} \rrbracket^- &:= \lambda M. u \ \overrightarrow{\langle N \rangle} M
\end{aligned}$$

Fig. 9. Optimized translation of Twelf signatures to λ Prolog programs

variables where δ contains the λ -bound variables currently in scope, while Δ contains the Π -quantified variables collected while walking through the type A .

Another, more direct, optimization is to reflect the LF types into types in the simply typed lambda calculus. Along with this optimization we can also use specialized predicates, rather than just **hastype**. For each LF type $u : K$ we will create a new atomic type **u-type** in *hohh*, as well as a new predicate **u** which has the type $\phi(K) \rightarrow \text{u-type} \rightarrow \text{o}$. We then use these to encode the signature in a more natural way. See Figure 9 for the new translation.

There are now two modes in which translation operates, the negative, $\llbracket \cdot \rrbracket^-$, which is essentially the same as before in that it does not check for strictness of bound variables, and the positive, $\llbracket \cdot \rrbracket^+$, which will only generate *hastype* formulas for variables which do not appear strictly. We do this to insure that the eliminations occur in situations in which it makes sense to think of the implication encoding an inhabitation check. We will write $\forall x. \llbracket B \rrbracket_{F,x}^+(M \ x)$ for $\forall x. \top \supset \llbracket B \rrbracket_{F,x}^+(M \ x)$ in future to simplify the generated signatures. These optimizations not only clean up the generated signature, but they also improve performance as we have limited the number of clauses which match the head of any given goal formula.

6 An illustration of the translation approach

We illustrate the use of the ideas described in the earlier sections by considering the **append** relation specified in Twelf by the signature in Figure 2. The Twelf query that we shall consider is the following that we previously saw in Section 2:

```
{x:nat} append (cons x nil) (cons z (cons x nil)) (L x).
```

This query asks for a substitution for **L** that yields an inhabited type and an object that is a corresponding inhabitant.

```

nat : nat-type -> o.
list : list-type -> o.
append : list-type -> list-type -> list-type -> append-type -> o.

nat z.
∀x. nat x ⊃ nat (s x).
list nil.
∀x.(nat x ⊃ ∀l. list l ⊃ list (cons x l)).

∀l. append nil l l (app-cons l).
∀x∀l1∀l2∀l3∀a. append l1 l2 l3 a ⊃
  append (cons x l1) l2 (cons x l3) (app-cons x l1 l2 l3 a).

```

Fig. 10. The Twelf specification of `append` translated into λ Prolog

Applying the optimized translation to the signature in Figure 2 yields the λ Prolog program shown in Figure 10. Further, the Twelf query of interest translates into the *hohh* goal formula

$$\forall x. \text{append} (\text{cons } x \text{ nil}) (\text{cons } z (\text{cons } x \text{ nil})) (L \ x) \ M.$$

The answer substitution for this goal in λ Prolog is

$$\begin{aligned}
L &= y \backslash \text{cons } y (\text{cons } z (\text{cons } y \text{ nil})), \\
M &= y \backslash \text{app-cons nil } (\text{cons } z (\text{cons } y \text{ nil})) \\
&\quad (\text{cons } z (\text{cons } y \text{ nil})) \ y \\
&\quad (\text{app-nil } (\text{cons } z (\text{cons } y \text{ nil})))
\end{aligned}$$

Applying the inverse translation described in Section 4 to this answer substitution yields the value for L and the proof term for the Twelf query that we saw in Section 2.

7 Conclusion

We have considered in this work an approach to implementing the logic programming treatment of LF specifications that is embodied in Twelf by using the Teyjus implementation of λ Prolog as a backend. Central to such an implementation is a meaning-preserving translation of Twelf specifications into λ Prolog programs. The basic structure of such a translation has previously been described by Snow et. al. [12]. Built into that translation is an optimization which takes advantage of statically available type information, quantified through a notion of strictness. In this work we have refined the notion of strictness to potentially enhance the usefulness of this optimization.

To actually use this approach in an implementation of Twelf, it is necessary to also provide a way of translating solutions found by Teyjus into LF terms that constitute answers to the query in LF syntax. Towards this end, we have presented an inverse encoding which describes how to map *hohh* terms back to LF objects in the context of the original Twelf specification.

The work by Snow et. al. deals only with terms which are closed, and so there had been no treatment for meta-variables which may appear in LF expressions. In order to capture the full scope of logic programming in Twelf, we extended the usual presentation of LF to allow for meta-variables in terms, and we provided a treatment for such variables in both the derivations and the translation. Although the proof showing the correctness of this translation is still incomplete, we have discussed an approach to developing such a proof that is based on relating the operational semantics of the two systems.

Acknowledgements

This work has been partially supported by the NSF Grant CCF-0917140. Opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. P. Brisset and O. Ridoux. The compilation of λ Prolog and its execution with MALI. Publication Interne 687, IRISA, 1992.
2. Alonzo Church. A formulation of the simple theory of types. *J. of Symbolic Logic*, 5:56–68, 1940.
3. Amy Felty and Dale Miller. Encoding a dependent-type λ -calculus in a logic programming language. In Mark Stickel, editor, *Proceedings of the 1990 Conference on Automated Deduction*, volume 449 of *LNAI*, pages 221–235. Springer, 1990.
4. Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
5. Robert Harper and Daniel R. Licata. Mechanizing metatheory in a logical framework. *Journal of Functional Programming*, 17(4–5):613–673, July 2007.
6. William A. Howard. The formulae-as-type notion of construction, 1969. In J. P. Seldin and R. Hindley, editors, *To H. B. Curry: Essays in Combinatory Logic, Lambda Calculus, and Formalism*, pages 479–490. Academic Press, New York, 1980.
7. Dale Miller and Gopalan Nadathur. *Programming with Higher-Order Logic*. Cambridge University Press, June 2012.
8. Gopalan Nadathur and Dale Miller. An Overview of λ Prolog. In *Fifth International Logic Programming Conference*, pages 810–827, Seattle, August 1988. MIT Press.
9. Frank Pfenning. Logic programming in the LF logical framework. In Gérard Huet and Gordon D. Plotkin, editors, *Logical Frameworks*, pages 149–181. Cambridge University Press, 1991.
10. Frank Pfenning and Carsten Schürmann. *Twelf User’s Guide*, 1.4 edition, December 2002.
11. Xiaochu Qi, Andrew Gacek, Steven Holte, Gopalan Nadathur, and Zach Snow. The Teyjus system – version 2, March 2008. <http://teyjus.cs.umn.edu/>.
12. Zachary Snow, David Baelde, and Gopalan Nadathur. A meta-programming approach to realizing dependently typed logic programming. In *ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP)*, pages 187–198, 2010.